

From the *IEEE Intelligent Systems Special Issue on Humanoid Robotics*, July/August 2000

Expanding Frontiers of Humanoid Robotics

Mark L. Swinson, DARPA

David J. Bruemmer, Strategic Analysis

Mobile robots pose a unique set of challenges to artificial intelligence researchers. Such challenges include issues of autonomy, uncertainty (both sensing and control), and reliability, which are all constrained by the discipline that the real world imposes. Planning, sensing, and acting must occur in concert and in context. That is, information processing must satisfy not only the constraints of logical correctness but also some assortment of crosscutting, physical constraints. Particularly interesting among these robots are humanoids, which assume an anthropomorphic (human-like) form.

A growing number of roboticists believe that the human form provides an excellent platform on which to enable interactive, real-world machine learning. Robots that can learn from natural, multimodal interactions with the environment might be able to accomplish tasks by means their designers did not explicitly implement and to adapt to the unanticipated circumstances in an unstructured environment. Ultimately, humanoids might prove to be the ideal robot design to interact with people. After all, humans tend to naturally interact with other human-like entities. Eventually, humans and humanoids might be able to cooperate in ways now imaginable only in science fiction. Humanoids might also provide a revolutionary way of studying cognitive science. As we review successes and failures in the field, we provide a contextual backdrop for understanding where humanoid research began, the dilemmas with which it currently struggles, and where it might take us in the future. We also discuss how these technological developments have and will continue to affect the ways in which we understand ourselves.

A brief history of "thinking" robots

In Plato's *Timaeus*, we find that the soul, before captivity within a human frame, knows no constraints while freely traversing the nonphysical realm. Yet, once inside the human body, the soul finds itself confounded by the inconsistency of the physical world, struggling to relate its prior knowledge of perfect, heavenly archetypes to the muddled reflections the senses perceive. Early attempts to build robots that could think and act like humans met a similar fate. Often derived in simulated environments, these agents possessed perfect, *a priori* knowledge of their virtual, archetypal worlds. Once embodied, these robots struggled to relate to a noisy and all too often inconsistent flow of data streaming in and out from a host of real-world sensors and actuators.

Understanding good, old-fashioned artificial intelligence. Instead of engineering effective, real-world behavior, classical AI emphasized computational intelligence. Researchers sought to implement rational thought processes and considered rational behavior to be an inevitable by-product. Researchers paid little regard to the correspondence problem as they constructed

Report Documentation Page

*Form Approved
OMB No. 0704-0188*

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE JUL 2000	2. REPORT TYPE	3. DATES COVERED 00-00-2000 to 00-00-2000		
4. TITLE AND SUBTITLE Expanding Frontiers of Humanoid Robotics		5a. CONTRACT NUMBER		
		5b. GRANT NUMBER		
		5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)		5d. PROJECT NUMBER		
		5e. TASK NUMBER		
		5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Defense Advanced Research Projects Agency, 3701 North Fairfax Drive, Arlington, VA, 22203		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)		
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited				
13. SUPPLEMENTARY NOTES				
14. ABSTRACT				
15. SUBJECT TERMS				
16. SECURITY CLASSIFICATION OF:		17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 9	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified			

increasingly complex and large knowledge-based systems to capture and process semantic information.

Researchers deemed symbolic representation paramount because it let agents operate on sophisticated human concepts and linguistically report on their action. As Donald Michie stated, "In AI-type learning, explainability is all."¹ The resulting emphasis on symbolic representation and planning profoundly affected robotics. Although these systems produced elaborate and elegant control architectures, the intelligence in these systems remained exclusively with the designer. The robots were merely automata executing static and often brittle programs.

Problems with hard-coded, top-down control. In their zeal to make robots think like humans, many researchers focused on high-level cognition and provided no mechanism for building control from the bottom up. Although intended to model humans, most systems did not, like humans, acquire their knowledge through interaction with the real world. When in the real world, these robots possessed little mastery over it. Even in the fortunate event that sensors could accurately connect internal archetypes to real-world objects, robots could only extend the knowledge thrust on them in rudimentary, systematic ways. Such robots carried out preconceived actions with no ability to react to unforeseen features of the environment or task.

Once a cause of great optimism, attempts to create human-like intelligence became a favored target for philosophical criticism. In 1979, Hubert Dreyfus argued that computer simulation assumes incorrectly that explicit rules can govern intellectual processes.² An ability to break rules, Dreyfus thought, better characterizes human intelligence. Rules allow only elementary capabilities and are routinely broken once we achieve true competence. He viewed this competence not merely as a new, more sophisticated set of rules but as the ability to serve principles that have not yet and might never become explicit. Another argument was that computer programs are inherently goal seeking and thus require the designer to know before hand exactly what behavior is desired (as in a chess match as opposed to a work of art).³ In contrast, humans are value-seeking—that is, we do not always begin with an end goal in mind but seek to bring implicit values to fruition, on the fly, through engagement in a creative or analytical process.

Although some of the ultimate conclusions were premature, these arguments aptly called attention to the fact that static programs, explicit rules, and knowledge bases drove robots estranged from the real world. As such, robots remained information-processing machines, applicable only to highly structured domains such as assembly lines. At best, those who claimed to be creating human-like intelligence were labeled positivists. At worst, they were considered delusional. Most roboticists forsook the goal of human-like cognition entirely and focused on creating functional, high-utility agents, using the lower animal world as a model (if they even needed models).

Toward a more robust, low-level knowledge. Realizing the limitations of hard-coded, externally derived solutions, many within the AI community decided to look to fields such as neuroscience, cognitive psychology, and biology for new insight. Before long, the multidisciplinary field of cognitive science drove home the notion that the planning and high-level cognition of which humans are consciously aware represents only the tip of a vast neurological iceberg.⁴ The mainstay of human action, researchers argued, derives from motor skills and implicit behavior encodings that lie beneath the level of conscious awareness.

Borrowing on this understanding, Philip Agre and David Chapman argued that robots should

likewise spend less time deliberating and more time responding to a world in constant flux.⁵ A new, behavior-based view of intelligence emerged that transferred the emphasis from intelligent processing to robust real-world action.

Neurobiology provided compelling evidence for a behavior-based approach with studies on the behavioral architecture of low-level animals. In one experiment, scientists severed the connection between a frog's spine and brain, effectively removing the possibility of centralized, high-level control. They then stimulated particular points along the spinal chord and found that much of the frog's behavior was encoded directly into the spine.⁶ For instance, stimulating one location prompted the frog to wipe its head, whereas another location encoded jumping behavior. This implicit, reactive control layer was what classical AI methods had ignored.

Bottom-up intelligence for more capable robots

For a new wave of roboticists, the question is how best to impart these primitive behaviors to robots. Attempts to directly hard-code such low-level behavior have proven either impossible or ineffectual. Instead, an increasing number of roboticists look to machine learning techniques, including artificial neural networks, genetic algorithms, and reinforcement learning. Neural networks provide a "supervised" learning approach where a designer trains a system's response to stimulation by adjusting weights between network nodes. Reinforcement learning provides an "unsupervised," learning-with-a-critic approach where systems can learn mappings from percepts to actions inductively through trial and error. Evolutionary methods begin with an initial pool of program elements and use genetic operators such as recombination and mutation to generate successive generations of increasingly better controllers.

Using these approaches and others, robots can learn by adjusting parameters, exploiting patterns, evolving rule sets, generating entire behaviors, devising new strategies, predicting environmental changes, recognizing the strategies of opponents, or exchanging knowledge with other robots. Such robots have the potential to acquire new knowledge at a variety of levels and to adapt existing knowledge to new purposes. Robots now learn to solve problems in ways that humans can scarcely understand. In fact, one side effect of these learning methods is systems that are anything but explainable. Careful design no longer suppresses emergent behavior but encourages it.

Applying learning approaches to humanoid robots

With the realization that the designer does not need to conceive solutions *a priori*, hope for building intelligent, human-like robots rekindled. By exploiting these learning techniques, roboticists have once again begun to tackle a variety of anthropomorphic capabilities. Many roboticists working with humanoids code learning mechanisms directly into their design environments and use them to hone existing behaviors, develop new behaviors, and string behaviors together. For instance, a designer can use a neural network to implicitly encode low-level motor control for an arm reaching behavior and then use reinforcement learning to train the humanoid when to reach and grasp. If the humanoid still struggles, the designer might, for instance, optimize behavior using a genetic algorithm to tweak parameters controlling rotational torque.

Although such methods have been invaluable, the devastating complexity of most humanoids has required specialization. The goal of human-like versatility has bowed to the goal of engineering specific human-like behaviors. The result has been humanoids that can exhibit impressive

functionality within a highly restricted domain or task. The next step is for an increasing number of capabilities to reside on general-purpose machines, engineered for all tasks because they are engineered for none in particular. Recent mechanical advances have produced humanoid bodies such as Robonaut (see "Robonaut: NASA's Space Humanoid," by Robert O. Ambrose, Hal Aldridge, R. Scott Askew, Robert R. Burridge, William Bluethmann, Myron Diftler, Chris Lovchik, Darby Magruder, and Fredrik Rehnmark in this issue) that represent an important step toward this goal.

Unfortunately, the software to enable such universal machines lags significantly. At first blush, the mechanical sophistication of a full-fledged humanoid body sounds like a devastating challenge to even the most robust learning technique. The more complex a humanoid body, the harder it is to place constraints necessary for productive learning. If we employ too few constraints, learning becomes intractable. On the other hand, too many constraints might curtail learning's ability to scale. Consequently, many of the most physically adept humanoid bodies tend to be driven by hard-coded behaviors or through a virtual reality human interface.

Ultimately, the conventional learning techniques we describe are perhaps most limited because they are tools human designers wield rather than self-directed capabilities of the robot. We submit that this might not need to be the case. Although robots will always require an initial program, this does not preclude them from indefinitely, willfully, and creatively building on it. After all, humans also begin with a program encoded in their DNA. The key is that in humans much of this genetic code is devoted not to mere behavior but to laying a foundation necessary for future development.

Automated development through real-world interaction

A growing number of humanoid researchers believe that this ability to appropriately seed development will make learning tractable for humanoids. The goal is no longer for robots to merely learn (acquire knowledge and skill in a particular area) but also to develop (enrich cognitive ability to learn and extend physical ability to apply learning). Truly autonomous humanoids must ultimately play some role as arbiters of their own development, must be able to channel and structure learning across layers of control. This will require generalized learning starting from the ground up and continuing throughout the humanoid's life, affecting what the robot is, rather than merely what the robot does.

Before we can transform a cognitive architecture into a developing mind, we must answer a host of difficult questions. How do we give humanoids the ability to impress their own meaning onto the world? How can humanoids direct their own development? How do we motivate this development? How much a priori skill and knowledge do we build in? Using what level of representation? What, if any, bounds should we impose?

Although these questions might never have definitive answers, an emerging learning approach provides a unique, functional balance of human input, self development, and real-world interaction. This approach, which we call *imitative learning*, lets the robot continuously learn through multimodal interactions with a human trainer and the environment. The robot does not simply process incoming information but actively responds to natural visual, auditory, and tactile stimulation. The robot can pose questions, ask for actions to be repeatedly demonstrated, and use emotional states to communicate frustration, exhaustion, or boredom to the human trainer.

Advocates of imitative learning see it as the cornerstone in a developmental foundation that can enable self-directed, future learning.

As you might expect, giving humanoids the ability to interact profitably with humans is not easy.

For imitative learning to succeed, robots must have some way of knowing which aspects of the environment to attend to and precisely which actions to reproduce. For instance, a robot should not imitate a cough or a scratch when a trainer shows it how to turn a crank. To guide robots through the process of imitative learning, we must give them the ability to recognize and respond to natural cues we give unconsciously through body language.

How far have we come?

This issue showcases a rich diversity of projects that use humanoid robots to model some subset of the physical, cognitive, emotional and social aspects of human body and experience. In "Social Constraints on Animate Vision," Cynthia Breazeal, Aaron Edsinger, Paul Fitzpatrick, Brian Scassellati, and Paulina Varchavskaya discuss an MIT project in which they are training a robot head called Kismet with eyebrows, eyelids, ears, and a mouth to discern and respond to social cues, such as nodding and eye contact, that are crucial in correctly guiding interaction. Bryan Adams, Cynthia Breazeal, Rodney Brooks, and Brian Scassellati discuss another robot platform, Cog, in "Humanoid Robots: A New Kind of Tool." They equipped Cog with a sophisticated visual system capable of saccades, smooth pursuit, vergence, and head and eye coordination through modeling of the human vestibulo-ocular reflex. Cog responds to visual stimulation, sounds, and the ways people move its body parts. By exploiting its ability to interact with humans, Cog can learn diverse behaviors including everything from playing with a slinky to using a hammer. Eventually, military commanders who might not know beforehand what tasks Cog will need to accomplish will be able to naturally and quickly task it.

Work with imitative learning also progresses at Michigan State University, where researchers are using *communicative learning* to iteratively hone behavior as the humanoid responds to verbal feedback from a human trainer.⁷ The foundational principle is that all human-derived forms of representation bias the system and inhibit learning's ability to scale. Instead, they wish the humanoid to build layers of control using as little built-in representation as possible. Rather than storing semantic information, the humanoid treats all stimulation as low-level vectors. Thus, the principles that let the robot process and learn from visual stimulation will apply equally well to other capabilities such as object manipulation.

Human-robot interaction plays a crucial role in the burgeoning market for intelligent service robots. Increasingly, robots that can serve as mobile, autonomous tour guides and information kiosks will grace public places. Sebastian Thrun, Jamie Schulte, and Chuck Rosenberg give an encouraging example in "Robots in Public Places: A Case Study." Their robot Minerva, a popular tour guide at the Smithsonian National Museum of American History, used a rich repertoire of interactive capabilities to attract people and guide them through the museum. Minerva's facial features and humanoid form greatly affected how people respond to it.

An ambitious effort at Vanderbilt University is working toward intelligent, task-general service robots that can aid the elderly and disabled. To deal with the complexity inherent to humanoid bodies and tasks, Kazuhiko Kawamura, R. Allen Peters II, D. Mitchell Wilkes, W. Anthony Alford, and Tamara E. Rogers ("Toward Interactive Human-Humanoid Teaming: Foundations and Achievements") designed their robot Intelligent Soft-Arm Control as a multiagent system that devotes a separate agent to each functional area. For instance, one agent deals with arm movement while another interacts with humans. Using *database associative memory*, ISAC can store and structure the knowledge it acquires. To mimic long-term memory, DBAM uses a spreading activation network to form associations between database records. To efficiently structure its memories, ISAC's Sensory EgoSphere processes incoming perceptual data according

to spatial and temporal significance.

In "A Neurobiological Perspective on the Design of Humanoid Robots and their Components," Simon Gistzer, Karen Moxon, Ilya Rybak and John Chapin provide a tour of recent neurobiological findings that continue to impact the world of humanoid robotics. The authors explain the process by which aspects of motor execution are encoded into a modular neural architecture within the spinal system. This architecture expedites motor control learning by constrainings the output sent to limbs and by hierarchically structuringes control primitives at varying levels. Maja Mataric's article ("Getting Humanoids to Move and Imitate") provides convincing evidence that roboticists can exploit this model by coding or training a set of basis behaviors on which developmental learning can build. Imitative learning is then a process of matching perceived behavior to an assemblage of these a priori primitives.

In "Using Humanoid Robots to Study Human Behavior," Chris Atkeson, Josh Hale, Mitsuo Kawato, Shinya Kotosaka, Frank Pollick, Marcia Riley, Stefan Schaal, Tomohiro Shibata, Gaurav Tevatia, Ales Ude and Sethu Vijayakumar discuss a collaborative, international endeavor that uses a 30 degree of freedom robot to emulate complex, full-body movement. For insight into human body movement, they use a unique motion capture system called a SenSuit which, when worn as an exoskeleton, allows researchers to record human movement trajectories for shoulders, elbows, wrists, hips, knees and ankles. This data helps to identify the underlying principles that constrain and optimize body movement. Ultimately, these principles will inform the way motion primitives are developed and used by humanoid designers. Currently, researchers have chosen to represent motion primitives using B-spline wavelets -- spikes in the kinematic graphs that characterize a specific joint movement. By providing an efficient way to specify and optimize multi-resolution motion trajectories, B-spline wavelets enable smooth, efficient movement.

In "Tracing Patterns and Attention: Humanoid Robot Cognition," Luiz-Marcos Garcia, Anotnio Oliveira, Roderic Grupen, David Wheeler, and Andrew Fagg use attentional mechanisms to focus a humanoid robot on visual areas of interest. On top of this capability, the authors have implemented a learning system that lets the robot autonomously recognize and categorize the environmental elements it extracts. They equip robots with perceptual clues such as sound, movement, color intensity, or human body language (pointing, gazing, and so on). For rich sensor modalities such as vision, perception is as much a process of excluding input as receiving it.

The articles in this special issue do not presume to exhaustively cover the realm of humanoid robotics. For example, in Japan, the electronics and automotive industries have played a key role in the resurgence of humanoids by developing robots capable of walking, climbing stairs, and even playing pianos. Although Japanese scientists have focused on the necessary mechatronics, they are also beginning to search for learning techniques that can scale indefinitely. At the University of Tokyo, researchers are using a learning methodology they call interactive teaching to give robots the ability to drive their own development. A robot uses Bayesian networks to map sensor evidence to behavior and then assigns each mapping a confidence rating. In the beginning stages, confidence ratings are low and the robot must frequently ask a human trainer for help deciding between competing actions. With practice, the robot requires less intervention from the human trainer until eventually it can autonomously complete a task. When the task changes, the robot can again ask for help.¹⁵

What does the future hold?

Although these projects are important steps in the right direction, functional results come slowly.

Like the human infants they model, developing humanoids are inefficient at most tasks and require intensive training. One implication of this research is that to create human-like adaptability and versatility, introducing an element of human frailty and inconsistency might be necessary.

As robots become increasingly pervasive, it remains to be seen whether humanoids can become crucial arbiters of this new world, able to favorably coexist with humans while exploiting the way we have structured our environment. Recent humanoid research has suggested that humanoid robots might one day perform surgery, build and maintain space stations, serve meals, or deliver packages throughout an office building. Moreover, researchers will task them naturally through gestures and speech. Nonetheless, there are still many who view humanoid research as a foolhardy, misdirected pursuit. Both inside and outside robotics, skeptics maintain that we could better spend money and time engineering targeted and arguably more affordable robotic solutions to fit specific needs.

Certainly, myriad tasks exist for which an ability to converse, learn, and interact is not necessary. For highly structured environments, factory automation robots are extremely adroit, efficient, and reliable. Even for tasks such as landmine detection, which might benefit from adaptation and autonomy, the robots do not necessarily need a human form or the ability to interact with humans. However, it takes little imagination to conceive the benefits of bringing highly capable humanoid agents to bear in scenarios spanning everything from firefighting or rescue operations to assisting the elderly and disabled. Moreover, such skepticism overlooks all that humanoid research can tell us about the way we think, learn, adapt, interact, develop, and evolve from an entity whose cognitive existence is not limited to a biologically constrained lifespan.

On the other hand, humanoids have much yet to prove. Will humanoid research propel robotics on to great heights, channeling ideas from diverse fields toward an ultimate goal? Or will the quest to model ourselves prove to be a stumbling block, or worse? We might be our best or worst models of intelligence. Although cognitive neuroscience will continue to contribute much to our self-understanding, we by no means fully appreciate the many internal processes that actually produce our intelligence.

Roboticist Rodney Brooks voiced similar sentiments, arguing that our view of how we think and act is tainted with subjectivity.⁹ We cannot wholly transcend our biased perspective. The best we can do is neutralize its effect by bringing humanoid bodies in line with our own. Most likely, we will never fully understand, much less recreate, everything that it means to be human. As the frontiers of our self-understanding expand, humanoid robots might simply follow (and at times propel) our continuously changing conception of what we are.

On one hand the human form can be viewed as an absurd and fragile vessel, ill-suited for any one task and redeemed only by the intelligence within. On the other hand, the human body provides us with a unique ability to learn and apply learning. Whereas dualistic thinking has often rendered the body little more than a tomb for the mind, it may be that, to the contrary, human-like intelligence requires a human-like body.

Humanoid robotics provides a unique forum in which to continue this age-old debate. Are humanoids destined to remain lumbering, overly complex, and ineffectual, or, like those they model, will they manage to grow into their ungainly form? This special issue attests that rather than hampering the application of AI, physical embodiment in the human form provides a necessary and useful grounding, letting humanoids surpass their original programming as they endeavor to communicate with their creators.

References

1. D. Michie, "Machine Learning in the Next Five Years," *Proc. Third European Working Session on Learning*, Pitman, London, 1988, pp. 107-122. [Back to Text](#) [Body](#)
2. H.L. Dreyfus, *What Computers Can't Do: The Limits of Artificial Intelligence*, Harper Colophon Books, New York, 1979. [Back to Text](#) [Body](#)
3. G.B. Kleindorfer and J.E. Martin, "The Iron Cage, Single Vision, and Newton's Sleep," *Research in Philosophy and Technology*, Vol. 3, 1983, pp.127-142. [Back to Text](#) [Body](#)
4. D.L. Schacter, C.Y.P Chiu, and K.N. Ochsner, "Implicit Memory: A Selective Review," *Ann. Rev. of Neuroscience*, Vol. 16, 1993, pp. 159-182. [Back to Text](#) [Body](#)
5. P.E. Agre and D. Chapman, "What are Plans For?" *Robotics and Autonomous Systems*, Vol. 6, 1990, pp. 17-34. [Back to Text](#) [Body](#)
6. S.F. Giszter, F.A. Mussa-Ivaldi, and E. Bizzi, "Convergent Force Fields Organized in the Frog's Spinal Cord," *J. Neuroscience*, Vol. 13, No. 2, 1993, pp. 467-491. [Back to Text](#) [Body](#)
7. J. Weng et al., "Developmental Robots: Theory, Method and Experimental Results," *Proc. Second Int'l Sym. on Humanoid Robots*, 1999, pp. 57-64. [Back to Text](#) [Body](#)
8. T. Inamura, M. Inaba, and H. Inoue, "Acquisition of Probabilistic Behavior Decision Model Based on the Interactive Teaching Method," *Proc. Ninth Int'l Conf. Advanced Robotics*, Japan Robot Association, Tokyo, 1999, pp. 523-528. [Back to Text](#) [Body](#)
9. R.A. Brooks, *Intelligence without Reason*, AI Memo No. 1,293, AI Laboratory, Massachusetts Inst. of Technology, Cambridge, Mass., 1991. [Back to Text](#) [Body](#)

Mark L. Swinson is the deputy director of DARPA's Information Technology Office and a US Army colonel. He is also the program manager for the Embedded Systems Program, the Mobile Autonomous Robot Software Program, and the Software for Distributed Robotics Program. His research interests include embedded software, distributed processing, domain-specific languages, and machine learning for robot programming. He has a BS in engineering from the US Military Academy at West Point, an MS in mechanical engineering from the University of Wisconsin, and a PhD in robot control systems from the University of Florida. He is a member of the American Society of Mechanical Engineers, the American Society for Engineering Education, the Association for Unmanned Vehicle Systems International, and the IEEE.

David J. Bruemmer is an engineering consultant at Strategic Analysis, Inc. where he provides technical support to robotics programs at DARPA. In addition, he is leading an effort to develop a portable, autonomous land-mine detection robot that can help combat land-mines, by exploiting recent sensor advances. He has a BA in computer science and religion from Swarthmore College. He is a member of Phi Beta Kappa, the Sigma Xi research society, the International Society for Adaptive Behavior and the American Association for Artificial Intelligence. His interests include autonomous agents, distributed robotics, adaptive behavior, and human-robot interaction. Specifically, he is interested in understanding the future implications of robotics technologies especially when applied to military and humanitarian purposes. Visit www.sainc.com/adaptiverobotics to learn more about machine learning, adaptive systems, and humanoid robotics.